

2.2 Statistical Approaches in Other Intrusion-Detection Systems

2.2.1 Sequences of Events in NIDES and Wisdom and Sense

The concept of a sequence of events is central to the Wisdom and Sense (W&S) system [7]. To set the stage for discussing how (W&S) handles sequences of events, we first discuss how both the statistical and rule-based components of NIDES handle events.

The NIDES statistical component can be considered as a tool for examining whether the "events" in short-term behavior are similar to the "events" in long-term behavior. Currently we define an "event" as a rather low-level concept, such as the execution of a particular command, logging on from a particular terminal, or accessing a particular file. A common feature of these events is that their occurrence can be determined from individual audit records. However, the statistical algorithms could be modified to handle higher-level events requiring information from multiple audit records. For example, we could define a compound event such as "logs onto a remote host as guest and then changes user name." Later, we discuss how higher-level events could be used to improve the detection capability of the NIDES statistical component. For now, we note that if higher-level events are added to the NIDES statistical component, these should not be restricted to those relatively few events that are recognized as being precursors or indicators of an attempt to compromise system security. Instead, the enumeration of suspicious events should be made in the rule-based component. As soon as an identified suspicious event is detected, the rule-based component can alert the security officer or take other preventive or monitoring actions. Again, the purpose of the statistical component is to detect a change in behavior rather than to identify the occurrence of a specific behavior.

The NIDES statistical component decides whether current behavior is anomalous based on the entire group of events that have occurred in a rapidly aged past (generally, the past few hundred audit records). Thus, the decision is based on a group of events, rather than any single event. NIDES raises an alarm when the most recent event pushes the short-term profile "over the limit" in the sense that the short-term profile becomes so different from the long term profile that the statistical component considers the short-term profile to be anomalous. Although the NIDES statistical component bases its decision on a group of events, it is not particularly influenced by the order in which those events occur (except that more recently observed events exert a greater influence on profile scoring than less recent events). Strictly speaking, therefore, the NIDES statistical component cannot be said to monitor ordered events.

The NIDES rule-based component does handle event sequences. In most cases with which we are familiar, the interest in sequences of events concerns very specific sequences that are associated with compromises of system security. The rule-based component can monitor the occurrence of different events, and when a sufficient por-

tion of the sequence has occurred, the rule-based component can notify the security officer that it is likely that the user in question is attempting to compromise system security. In accordance with our general design philosophy, if a sequence has been specifically identified because it has been associated with compromised security, it belongs in the rule-based component. If, on the other hand, we are examining sequences so that we can better understand normal behavior, and therefore contrast short-term behavior (i.e., recent behavior) with long-term behavior (i.e., historical behavior), the proper place to examine the sequence would be in the statistical component.

Many rule-based systems (including the NIDES rule-based component) monitor sequences of events. The only statistical intrusion-detection system with which we are familiar that uses sequences is Wisdom and Sense, and the types of sequences considered by W&S are fairly elementary. The W&S system builds its own sequences and must consequently deal with a huge number of possible sequences. It does so by building a massive tree structure (the branches of which are the sequences) and then pruning the tree so that it contains only the most frequently used sequences. However, to reduce the overhead required to build the tree structures, the W&S system usually does not build tree structures for each individual user (instead, building a single tree for all users) and the tree structure involves data from only the past few days. Recent behavior of the user is compared to this tree structure in the following sense — whenever the user has executed a sequence of n events, the first $n - 1$ of which are in the tree, the n th event is evaluated to determine whether or not it is also in the tree. If it is in the tree, it typically nudges W&S toward concluding that the user is behaving normally (particularly if the last node representing that event has a high conditional probability of occurrence given the first $n - 1$ nodes). If it is not in the tree (or if the conditional probability of the last node is low), the last event nudges W&S towards concluding that the user is behaving abnormally. Typically, sequences in W&S are less than four to six events long. In addition, typical events are fairly elementary in nature. For example, an event might be a user name, an action, an object for an action, or a time of occurrence for an action. Thus, the mega-event that user A logged in from terminal A at 8:20 a.m. would constitute a four-event sequence represented in a four-node tree structure.

Owing to its complexity and the relative lack of documentation concerning W&S, we do not know whether the W&S system of considering sequences of elementary events and then evaluating conditional probabilities for the last node of those sequences is more or less powerful than the NIDES statistical component. We have not chosen to implement the W&S approach, primarily because the W&S system tends to profile a large group of users over very few days, whereas we wanted to base our statistical comparisons on individual profiles of longer-term behavior. However, the W&S approach does have merit, and its approach to handling sequences could presumably be incorporated into NIDES by building hooks into W&S's code.

Because of the potential value of sequences of events, we are also interested in determining how sequences could be incorporated into the NIDES statistical compo-

ment. The easiest way of doing this would be to define "mega-events." For example, the higher-level or compound event considered earlier (i.e., log onto remote host as guest and then change user name) is a simple sequence of events. Mega-events could be incorporated into the NIDES statistical component without a change in the current methodology. However, the definition of those events, and the development of software to recognize those events, would be a substantial research task.

Another method for embedding sequences into the statistical component would be to rewrite the code so that the statistical component builds its own sequences. For example, the statistical component could identify and track all three-event sequences and monitor their likelihood of occurrence. The difficulty with this approach is that the number of different sequences being monitored could become excessively large. For example, even if there were only 20 basic events, the number of possible sequences of three events would be $20 \times 20 \times 20 = 8000$. Unless the user's behavior were concentrated on a very small portion of these events, or the events could be collapsed into a smaller number, the sequences would not be useful for examining abnormality (because each individual sequence would have an extremely small probability of normal occurrence). Methods for collapsing (or clustering) behaviors into a reasonable number of mega-events or sequences would be necessary. We believe that it might be possible to develop a reasonable collapsing strategy, and that the resultant methodology would be more powerful than the one used in W&S, because it would not require the asymmetry inherent in the W&S approach (wherein most measures serve as "conditions" and only one measure at a time can act as the "outcome") and would more fully exploit the "multidimensionality" of the data.

2.2.2 Haystack's Statistical Approach

In many respects, the Haystack statistical approach appears to be similar to the NIDES statistical approach, although there are some quite distinct differences, and we discuss both these similarities and differences here. We caution that we have not been able to obtain complete documentation concerning Haystack [11], and therefore our conclusions regarding this system are preliminary. Steve Smaha, the Haystack system developer, also recommended that we note that the Haystack system is a number of years old.

In Haystack, as in NIDES, individual measures (called features) are monitored and compared to historical behavior to determine their abnormality. However, in Haystack this assessment appears to take place at the conclusion of the session, rather than in real time or across sessions. In addition, it appears that only "counting" measures (such as amount of I/O, amount of CPU, or number of files) are supported. (NIDES also supports "categorical" measures such as name of terminal, names of files accessed, names of directories changed, and names of commands executed.) For each measure, the Haystack system determines a range of values that are considered "normal." For example, if the measure is number of files accessed, the normal behavior for a user

Rationale

21

consists of a range (such as 0 to 34 files) containing at least 90% of all observed values for number of files accessed in the sessions belonging to the particular security group to which the user belongs.

In contrast to NIDES, the Haystack statistical system contains a set of six generic types of computer abuse (attempted break-ins by unauthorized users, masquerade attack, penetration of the security control system, leakage, denial of service, and malicious use). For each type of computer abuse there is a set of weights (from 0 to 9) indicating the extent to which each measure is related to that type of computer abuse. For example, for denial of service, the measure "number of pages printed" might be given a weight of 9, whereas a measure "password errors" might be given a lesser weight. When Haystack analyzes a session, each session feature outside of the predefined range causes a weight for that feature to be added to the session's score for that computer abuse. Each of the features is assumed to be independent of the other features. Under this assumption, it is possible to calculate the probability distribution of the computer abuse score (called a suspicion quotient) and alert the security officer when that score is too large.

Haystack also contains a simple "trend" test to determine whether the counting features are trending to larger or smaller values. This test (the Mann-Whitney procedure) is applied to contrast the values of the feature in the most recent 5 sessions with the values of the feature in the most recent 20 sessions. Basically, this test looks for changes in the median values. This test would not be applicable to categorical measures, nor to more complicated types of changes in counting distributions (for example, a trend towards both larger and smaller values but not a change in the median value).

Although we have never conducted any tests with the Haystack system, we believe that the NIDES methodology for handling individual measures is preferable to the Haystack methodology. For counting measures the difference in effectiveness is probably marginal, except when the measures are bimodal (such as would occur if a user engaged in two primary types of activity, one with low CPU usage per session and one with high CPU usage per session), in which case the SRI approach should yield better results. However, we believe that the most powerful measures will be the categorical measures, which the Haystack system does not support. In addition, the NIDES system supports nearly real time intrusion detection, whereas the Haystack system supports only retrospective analysis of entire sessions.

The Haystack system combines functions of the NIDES statistical component (i.e., identifying abnormal behavior) with functions of the planned NIDES resolver (i.e., drawing conclusions as to whether a particular behavior is suggestive of a type of computer abuse). Rather than combine these two distinct functions, we have chosen to separate them, believing that by doing so each will be more efficacious. For example, the planned resolver functions, to include model-based definitions of suspicious behaviors, could potentially be far more powerful than Haystack's relatively simple approach.

The degree of success of Haystack's approach for drawing conclusions depends on the validity of the weights used to obtain scores for computer abuse types. Unfortunately, we have not been able to obtain those weights (or even a complete listing of the measures used). We would be interested in examining the Haystack weighting system to assess its sensitivity in detecting simulated intrusions. (In addition, we note that in the absence of further development of model-based reasoning in the resolver, the Haystack weighting system could be implemented fairly readily within the NIDES resolver, while maintaining real-time intrusion detection and the use of categorical measures.)

2.3 Alternative General Statistical Approaches

2.3.1 Pattern Recognition

Pattern recognition [2] is a fairly well defined set of procedures when applied to the problem of identifying physical objects in two- or three-dimensional visual representations. Thus, when we discuss "pattern recognition" with respect to the problem of having a robot navigate a room filled with objects or the problem of identifying objects in a high-altitude photograph, it is fairly clear what set of procedures we are considering.

In other realms of statistics, pattern recognition does not refer to a well-defined set of procedures. A very wide range of statistical procedures has been used, in some sense or another, to find "patterns" in the data. As an extreme example, even regression analysis (which is not usually considered to be a form of pattern recognition) bears many characteristics of pattern recognition, since it is an optimal procedure for determining how well the patterns of values in the independent variables match (or fit) the pattern of values in the dependent variable. Thus, one of the main difficulties in asking "Why didn't NIDES use pattern recognition as its statistical approach" is that so many different statistical procedures qualify as pattern recognition approaches.

In fact, we could easily consider the NIDES statistical approach to be one of pattern recognition. The NIDES statistical component builds a long-term profile (or pattern) of user behavior and then compares this pattern to the short-term profile (or pattern) of user behavior. Thus, the NIDES statistical component is not only constructing patterns, but is also comparing patterns.

Probably, the statistical procedures that are most closely associated with pattern recognition are cluster analysis [1] and factor analysis [8]. Both of these procedures are essentially exploratory in nature. We use cluster analysis to find clusters in multivariate space (typically higher than three dimensions) and factor analysis to reduce the dimensionality (i.e., number) of outcome variables.

With respect to factor analysis, it is difficult to conceive how such an approach would be applied to intrusion detection. Presumably, each event would be characterized according to measures such as the amount of CPU used, the number of files

Rationale

23

accessed, and the time of day of the activity) and then this data would be processed to find fewer "factors" that explain the correlation among these measures. However, we do not believe that the results of such a factor analysis would be interpretable, and would probably only confuse the security officer.

It might be possible to use cluster analysis for anomaly detection. In this case, we would characterize events according to specific measures (such as amount of CPU used, the numbers of files accessed, whether the user accessed some other user's directory, and the time of day of the activity). We would then locate events gathered over a fairly long period of time (say weeks or months) in n -space (where n is the number of measures used in the event characterization) and form clusters. A probability would be assigned to each cluster corresponding to the number of events in that cluster (relative to the total number of events). The events in short-term behavior would also be identified, where we might define short-term behavior as the last 20 events. A running score would be computed for each user. That score would increase if the last event engaged in by the user did not belong to any previously identified cluster or was identified with a cluster with low probability of occurrence. Similarly, that score would decrease if the last event engaged in by the user belonged to a previously identified cluster with a high probability of occurrence. When the score became sufficiently large, the security officer would be notified that the user was engaging in anomalous behavior.

The use of cluster analysis in the way described above would be compatible with the current NIDES statistical approach. In essence, we currently track the probability of each event by itself and increase or decrease the anomaly score function on the basis of that individual probability. The cluster analysis approach would allow us to cluster events into larger collections of events and then assign probabilities to that cluster (which would be the sum of the probabilities of the individual events within that cluster).

The difficulty in applying cluster analysis is that it is more an art form than a set of definitive procedures. We are aware of and have used a wide variety of cluster analysis procedures on other projects. Unfortunately, a large amount of a statistician's time and attention is necessary to make the clusters meaningful. For example, the definition of the distance metric plays a critical role in determining which events would cluster together and in determining whether we get clusters that look more like "balls" or "strings" of events. We believe that it would be a fairly substantial research effort to identify reasonable metrics and clustering procedures. Worse, the clustering would have to be done repeatedly (say every few nights) on hundreds or thousands of individual users, because it is doubtful that clusters would be generalizable across systems or user communities. This would require storage of immense amounts of data (for example, the last 10,000 or so events for each user of the system) and might involve exorbitantly long processing times since cluster analyses are time-consuming. Unfortunately, the cluster analysis would probably frequently fail. Cluster analysis procedures tend to be very sensitive to the identification of

the proper "seeds" for the clusters, and often it takes a trained analyst to identify which seeds to use, based on the results of an unsuccessful clustering attempt. Thus, although cluster analysis is a very useful exploratory tool we do not believe that it would be appropriate for use in an anomaly detection system.

2.3.2 Discriminant or Classification Analysis

Classification analysis relies on previous examples of behavior to develop a set of rules upon which current behavior can be classified. For example, if we had historical examples of "normal" and "abnormal" behavior for a user, we could use those examples to classify new behavior as either normal or abnormal. Two commonly used procedures are discriminant analysis and classification tree analysis [6].

In our experience, classification analysis is an extremely efficient and powerful methodology for classifying current behavior as normal or abnormal. The principal problem in using this analysis procedure is in finding the examples of abnormal behavior. We simply do not have a sufficient number of examples of behavior that compromise system security. Nor do we trust the few examples that we have to be exhaustive, in the sense that they cover all types of behavior that could compromise system security. Even if we had exhaustive examples of known behavior that compromised system security, we would have to be concerned about the extent to which unknown or new behavior that compromised system security would resemble previously identified behavior that compromised system security.

Because actual examples of behavior that compromises system security are difficult to come by, analysts sometimes use the normal behavior of other users to characterize abnormal behavior for the host user. The first prototype for IDES [4] used both discriminant analysis and classification tree analysis in this way, and we showed that it was relatively easy to distinguish between the normal behavior of different users. We did not choose to continue the development of a classification approach because we had doubts concerning the extent to which behavior intended to compromise system security would resemble the normal behavior of other users. However, this approach may still have some value in identifying masqueraders, and we have performed some preliminary evaluations showing that there may be merit in adding a discriminant analysis into NIDES.

It is possible to heuristically relate the current NIDES statistical approach to a classification approach. Basically, in the current NIDES statistical approach we calculate the absolute probability of each type of event (as determined by examining the historical behavior of the host user) and then score recent behavior on the basis of that absolute probability. Thus, if the absolute probability of an event (as measured by the past history of events for that user) is small (say on the order of 0.01%), and the user has recently engaged in that event, we nudge the anomaly score upwards towards a conclusion of "anomaly." On the other hand, in classification analysis we calculate the relative probability for each type of event, and then score recent behav-

ior on the basis of that relative probability. For example, if the user engages only in a particular type of event with a probability of 0.01%, but the comparison population (i.e., all other users or intruders) engages in that same type of event with a probability of 0.0001%, then the occurrence of that event is 100 times more likely for the host than for the comparison population. Therefore, if that event occurs in recent behavior it is evidence in favor of normal behavior rather than abnormal behavior. Thus, it is possible (although not highly likely) that classification analysis and the current NIDES analysis procedures could lead to opposite conclusions. The critical role played by relative probabilities means that the definition of the comparison population and the inclusiveness and representativeness of the examples of the behavior of the comparison population are critical to the success of classification analysis.

2.3.3 Markovian Transition Analysis

In Markovian transition analysis [5], we are typically concerned with the definition of different states and the probabilities of transitions from one state to another.

In intrusion analysis, Markovian transition analysis might be used to anticipate to which state a user would move. This would be important if the states were associated with compromised security and it was necessary to take preventive action prior to a particular state's being reached. However, in this event it is necessary to calculate this probability only once. Thereafter, a rule could be entered into a rule-based component that simply states that if state A occurs the probability that the next state will be associated with compromised security is higher than the allowable threshold, and therefore preventive action should occur.

Another way to use Markovian transition analysis would be to characterize the normal behavior of a user on the basis of which states he or she moves between and with what probabilities. This approach would be absolutely consistent with the current NIDES approach, where an "event" would be the transition between two particular states. NIDES would track the probabilities of these events (i.e., the transition probabilities between states) and would score a user on the basis of whether or not the transitions were typical for the host user. The only additional effort needed to implement this approach is that involved in defining the states that we want to track movement between, and writing the software to detect that a state has been achieved. This is not an easy exercise; we note that for N states, there are N^2 transition probabilities, and therefore the transition matrix approach can quickly become unmanageable.

2.3.4 Bayesian Decision Analysis

The NIDES statistical team has recently implemented one of the most complicated and state-of-the-art Bayesian decision analysis systems in use today, and is therefore familiar with the advanced theory of Bayesian decision analysis as well as the imple-

mentation issues involved.[10] We believe that this analysis procedure (or a similar procedure such as evidential reasoning) should play an important role in NIDES. However, we believe that the proper role is in the resolver, which will be able to accept information from both the statistical and rule-based components and then draw conclusions with respect to the likelihood that system security is being compromised and the seriousness and type of the potential breach.

To some extent, a Bayesian decision analysis can be considered the next step beyond the traditional rule-based system. In the traditional rule based system, rules either fire or they do not. That is, each state represented by the rule has a probability of either 0 or 1. There have been a number of attempts to integrate probabilities into rule-based systems, so that we can infer the probability that a state is true, rather than definitely judging it true or false. Most of these systems work to some extent, but sooner or later run afoul of the rules of probability theory. (For example, these systems rarely handle contradictory information very well.) Bayesian decision analysis systems can be considered to be rule based systems that do obey the rules of probability theory (although the rules are so cleverly disguised in the relationships among nodes in the probability network that it is often difficult to realize that there are even rules present).

To be truly successful, a Bayesian decision analysis system must incorporate a substantial amount of knowledge concerning the different types of attacks that can be used to compromise system security as well as the conditional probabilities that various well-defined events will occur given that those attacks are in progress. We believe that the intrusion-detection community is only at the first stages of trying to assemble this type of knowledge.

Because the Bayesian decision analysis system approach relies very heavily on our current knowledge of methods for compromising security, we do not believe that it is a substitute for the current NIDES statistical component, which attempts to identify any behavior that is not consistent with past behavior.

Chapter 3

Evaluation of Statistical Approaches for Appropriateness to Intrusion Detection

In the process of developing the NIDES statistical component, and in reading the literature on intrusion detection, we had the occasion to consider a number of different statistical approaches that might have been used. Although we did not develop a set of formal evaluation criteria, it quickly became evident to us that most statistical approaches were not suitable for intrusion detection. For example, the method that we developed for NIDES, while it bears some similarity to chi-square analysis, has undergone substantial modification and is much different from that approach as found in the statistical literature.

Here, we present a set of evaluation criteria that should be considered when determining how likely it is that a proposed statistical approach will work in a real-time intrusion-detection environment.

3.1 Criterion 1. Does the Method Depend Upon Distributional Assumptions?

Many statistical approaches require that the data fit a normal or other distribution (such as gamma, chi-square, beta). Unfortunately, these assumptions are rarely true. This can render the statistical approach unsuitable, unless one of two circumstances occurs:

- The approach transforms the data in such a way as to force these distributions to be true. For example, the average of a large number of observations would probably be approximately normally distributed even if the individual observations were not. Another example is in NIDES, where we force the measures to

follow what is essentially a chi-square distribution by mapping the empirical cumulative distribution function of the measure into a "half-normal" distribution and then squaring this variable.

- The statistical approach is robust to deviations from the distributional assumption. For example, nonparametric approaches have few or no distribution assumptions.

We note that it is typically not sufficient that the data for which a distribution is assumed pass a goodness-of-fit test (which typically verifies only that the bulk of the probability distribution has the required shape), because we are typically interested in the extreme tails of the distribution of the test statistic.

The NIDES statistical component was developed so that it did not depend on distributional assumptions. We felt that this was necessary because the system was intended to be applied across a wide variety of computer systems, types of users, and types of measures. Our reluctance to assume any distribution constrains the types of statistical procedures that could be used. For example, if the measures could be assumed to follow a true multivariate normal distribution, then there are optimal statistical procedures that could have been applied.

3.2 Criterion 2. Does the Method Assume Independence?

Data from audit streams are notoriously interdependent. For example, once a file has been opened and read, it is highly likely that the following audit records will contain subsequent reads on that file. Furthermore, the correlation that we tend to encounter is not simple serial correlation. For example, if two files are open, then accesses to these files are likely to alternate over a time span of audit data (as would be the case if one file was open for reading and another for writing).

For a statistical approach to be suitable for intrusion detection it must not depend upon the independence of the audit trail records. This is a very strong condition and is one of the primary reasons why statistical approaches fail to be appropriate. The NIDES statistical component has a feature that accommodates the lack of independence of successive audit records without requiring that the type of dependence be modeled. This feature is essentially the empirical observation of the distribution of the evaluation statistic obtained by observing its values over time, rather than postulating its distribution based on the observation of the distribution of its component parts and assuming the independence of those parts.

3.3 Criterion 3. Does the Method Accommodate Categorical Data?

In our experience, the most valuable data in the audit trail are the categorical data, such as names of directories accessed, names of files accessed, commands used, and terminals used. Furthermore, there are sometimes hundreds (or even thousands) of categories for a single measure. Unfortunately, many statistical approaches assume the existence of continuous data and are not capable of handling categorical data well. Such approaches are less likely to extract sufficient information out of the data stream to be useful (i.e., to have a high probability of detecting misuse or an intrusion). Much of the recent work in NIDES has gone into making our approach very friendly to categorical data (in fact, we even transform continuous measures into categorical measures).

3.4 Criterion 4. Does the Method Allow Real-Time Evaluation of Audit Records?

During a typical day, tens or hundreds of thousands of audit records may be generated, and these records will tend to be heavily concentrated in about 10 business hours. If the objective is real-time audit trail assessment, then the statistical component will similarly have to implement tens or hundreds of thousands of assessments each day of whether an intrusion or misuse has occurred. This imposes considerable constraints on the statistical procedure — it must be sufficiently efficient to execute an assessment in a fraction of a second. (NIDES is currently capable of processing approximately 30 to 45 audit records per second.) Although many statistical procedures are capable of this (e.g., discriminant analysis, neural networks), other procedures may not be capable. In particular, there may be a problem if the proposed statistical approach is using a sliding window of data (say the last 100 audit records), and whenever a new audit record is added and one subtracted from the historical queue, the entire statistic needs to be recomputed, rather than merely being given a simple recursive adjustment. Many of the refinements to the NIDES statistical component were implemented in order to speed evaluation, including the use of exponentially weighted “sufficient” statistics.

3.5 Criterion 5. Does the Method Allow for Profiles for Individual Users?

Individual profiles may vary substantially from person to person. In order to achieve the maximum discriminatory power, we believe that in general, it is necessary to develop individual profiles. In developing NIDES, we set as a requirement that each

user have his or her own profile. This requirement affects both the time necessary to update profiles (since there are many profiles to update) and memory requirements. Profiles must be capable of being stored in a reasonable amount of space, or an excessive amount of time will be spent swapping profiles in and out of memory. Currently, NIDES requires about 100kb per profile in the Sun-Unix environment, which has a rich set of measures available. We suspect that the NIDES profiles could be stored in less space when using audit record data from other computer systems.

In some environments job classifications may be tightly prescribed, and individuals in the same job classification may use the computer in essentially the same way. In this case it may be possible to develop useful group profiles, and individual profiles may not be necessary. A statistical approach used in such an environment would not need to allow for individual user profiles, because group profiles would be sufficient. However, because NIDES was developed as a general-purpose intrusion-detection system, we developed a statistical component that could efficiently process individual profiles.

3.6 Criterion 6. Does the Method Allow for Profiles to Periodically Update without Human Intervention?

User profiles are constantly changing as the legitimate work of the user changes. Therefore, it is necessary that the statistical procedure be capable of updating user profiles at least once daily. Because there may be thousands of user profiles, efficiency in updating is important.

If we assume that updating should occur once per day, and that during the 10 business hours of the day the statistical system will be too busy processing audit records to update profiles, then the profiles will need to be updated during the remaining 14 hours (i.e., 50,000 seconds) of the day. If there are 2000 user profiles (i.e., a decently large system), each profile will need to be updated within 25 seconds. (With fewer users the profile updating time can take longer, and with more users the profile updating time will need to be less.) If we put in a safety margin of 50%, then profiles for a 2000 user system should each update within 12.5 seconds. While most statistical algorithms are capable of updating in this time frame, some may require more than this amount of time. The NIDES statistical component can update all profiles for an individual user in a second or so, and consequently if we initiate updating at midnight, we are typically finished within tens of minutes later.

It is also important that the updating procedure not require the guidance of a trained statistician, but rather be completely automated. Some statistical procedures, such as cluster analysis, typically require the intervention of a statistician to obtain decent results. A statistician will not typically be available when updating is required, and even if he or she were available, human intervention would slow the updating

process unacceptably. The NIDES statistical component does not require human intervention to update profiles.

3.7 Criterion 7. Does the Method Allow for Multivariate Statistical Inference?

It is our belief that much of the ability to detect interactions rests in the interaction of different aspects of behavior. For example, a set of particular commands may be associated with a particular host computer or with action on a particular set of files. The statistical approach should therefore be able to extract additional information about the normality of different actions taking place essentially simultaneously (e.g., invoking commands and accessing files).

The use of highly dimensional multivariate data involves many difficulties that affect the ability of the statistical approach to satisfy other criteria. For example, the multivariate cross product of only three measures, with only 100 categories each, results in a matrix of 1,000,000 possibilities. Even though the matrix would probably be sparsely populated, storing and processing such a matrix would tend to make the statistical algorithm too slow and the storage requirements too large. In addition, development of stable probabilities for the matrix cells would require long training periods — perhaps sufficiently long that behavior would change before training was completed.

The NIDES approach makes limited use of multivariate data, and we consider this to be one of the statistical component's primary limitations. However, we have developed the theoretical approach necessary to incorporate multivariate data (at least to some extent), and could incorporate this approach in future versions of NIDES.

3.8 Criterion 8. Does the Method Require the Existence of Defined Jobs or Sessions?

Computing used to be defined in terms of batch jobs. Now most processing takes place in continuous sessions, which may last for many hours. (In the extreme case, the user may never log off.) The statistical approach should not require that there be "jobs" or "sessions" or their equivalents, but rather should make assessments about data from a continuous stream of audit records. The lack of a natural time unit of analysis (that is, job or session) is directly responsible for the use in NIDES of exponentially weighted statistics.

3.9 Criterion 9. Does the Method Require the Existence of Simulated or Actual Intrusions?

Many statistical classification methods such as discriminant analysis or some neural networks require that training sets exist and the audit data in the training set be preclassified. For example, one portion of the training set might consist of audit data from users who are known not to be misusing the computer system and who own accounts that have not been subjected to intrusions. The other portion of the training set might consist of audit data gathered during simulated or actual misuse or intrusions. By comparing the "normal" and "intrusive" audit records in the training set, the statistical method can determine what features in the audit data best classify the data as normal or intrusive.

Unfortunately, data sets of simulated or actual intrusions are not widely available. (If such data sets had been available, we would probably have included a statistical classification method in NIDES.) A researcher who proposes to use a classification approach that requires the existence of simulated or actual intrusions should have a good idea where such data would come from or how to generate it. In addition, the ability of the statistical method to detect intrusions will generally be limited by the type of intrusions that it is trained on. That is, it will be able to detect new types of intrusions only to the extent to which they share features with the types of intrusions in the training set. Therefore, the generality and completeness of the simulated or actual intrusions in the training set should be assessed.

3.10 Criterion 10. Does the Method Develop Its Assessment Based on Differences between Users?

Because training sets of simulated or actual intrusions are not typically available, the researcher might propose to apply the statistical classification method to a training set consisting only of normal users. For example, to develop an algorithm for identifying abnormalities in User A, one portion of the training set might consist of user A's audit data and the other portion consist of a mixture of audit data from other users (which might be identified either as "non-A" or by their individual user IDs). The statistical method would then be trained to classify a new audit record either as A or as non-A.

Statistical approaches that distinguish normal users from one another might also be capable of identifying misuse or intrusive activity in an audit stream. We believe that such methods are worth investigating. However, it is not clear that such methods would actually be able to detect unauthorized activity. For example, for there to be a high probability that unauthorized activity in User A's account would be identified

as non-A, that unauthorized activity must much more closely resemble the typical activity of other users than the activity of user A. In the event that the intrusive activity does not closely match either user A's activity or the activity of other users, statistical approaches such as discriminant analysis might fail to classify the activity as not belonging to A.

The NIDES methodology continuously compares a user's recent activity with his or her long-term profile. It is in this sense self-referential and independent of the activity of other users on the system. As such, it makes no assumption as to the nature of the difference between legitimate and unauthorized activity.

3.11 Criterion 11. Does the Method Require Assumptions about the Distribution of Intrusive Behavior?

Because data on actual and simulated intrusions are scarce, and a number of statistical procedures have been proven to have optimality properties when comparing two known distributions, it is tempting to make assumptions concerning the distribution of intrusive behavior. These assumptions are typically of one of two types. The first is an assumption that intrusive behavior always tends towards larger (or sometimes smaller) values of a particular measure of resource consumption (e.g., CPU time, I/O, file accesses, "finger" commands). The second is an assumption that intrusive behavior is uniformly distributed across all possible behaviors.

Although these assumptions allow "optimal" tests to be derived, unless the correctness of these assumptions is verified, the optimality of these tests must be severely discounted. In fact, most "reasonable" statistical tests can be shown to be optimal or near optimal versus some assumed distribution of intrusive behavior. This can be done by finding the set of behaviors that would result in the declaration of an anomaly and then assuming that intruder behavior is concentrated on those outcomes. (Unfortunately, the math involved in making such assessments is quite complicated.)

The NIDES statistical algorithms were not postulated with a specific distribution of intrusive behavior in mind, nor have we reverse-engineered the set of intrusive behaviors for which our statistic is optimal. However, the form of our test statistic (e.g., similar to a chi-square statistic) suggests that we spread our detection power over many possible intrusive behaviors and so would do particularly well versus a flat or uniform distribution for intrusive behavior.

3.12 Criterion 12. Does the Method Provide the Security Officer with Adequate Information to Conduct an Inquiry?

By their very nature, statistical approaches to anomaly detection will raise false positive alerts. That is, users do change their behavior over time — and sometimes abruptly so. Statistical algorithms will detect these changes and issue alerts even though no misuse activity is actually transpiring. Because the vast majority (i.e., 99% or more) of activity on a typical computer system will be legitimate, the number of false positive alerts will tend to outnumber the true positive alerts. A good deal of the system security officer's time will therefore be devoted to investigating false positive alerts. The statistical procedure must be capable of providing the security officer with information that can be used to identify the reason why the alert was issued. This will tend to greatly reduce the time involved in the security officer's determination of whether misuse has actually occurred.

The requirement for "understandability" may restrict the type of statistical procedures that can be applied. For example, factor analytic procedures tend to transform the space of measures in complicated ways, and it can become quite difficult to unscramble the statistical results in such a way that they can be explained to the security officer.

Because of the need for understandability, the NIDES statistical component uses the concept of long- and short-term profiles, which are probability distributions for categorical measures. Although the presentation format for these probability distributions still requires improvement, the concept of frequency distributions is easily grasped. For example, the security officer can see how frequently different commands have been issued in the past and which commands have been more recently issued. It is relatively easy then to identify those commands that are now being issued, but have not been issued in the past (or were issued much less frequently in the past) as well as those commands that were frequently issued in the past but are now conspicuously absent. The use of relatively simple individual measures (e.g., commands used, file names, hosts used, number of access denials) makes the results more comprehensible.

Chapter 4

Answers to Specific Questions about Features of the NIDES Statistical Algorithm

The NIDES statistical approach has evolved over many years. This evolution has been motivated by a mixture of theoretical considerations, practical considerations (dealing principally with issues of computing speed), and practical experience in exercising the algorithm. Here, we identify some of the major decisions that shaped the statistical approach and describe some of the specific features of the algorithm. We have adopted a question-and-answer format, where the questions deal with different features of the NIDES statistical approach, and the answers offer insight into why those features were incorporated. Occasionally, we group a set of questions that have closely related answers. We have also formulated a number of the questions (particularly the questions that appear early in this section) so that they illuminate the relationship between NIDES and the approach of Helman et al.[9]

4.1 Q: Is the NIDES Statistical Approach Based on an Assumed Intruder Behavior?

A decision made early in the NIDES development process is that we should not assume any knowledge of how intruders would compromise system security (or how a user would alter his or her behavior to compromise system security). All knowledge of this nature was to be incorporated into the expert system. We made this decision for four principal reasons:

- The expert system is typically more accurate in detecting known intrusion approaches.

- There were insufficient data on intrusion attempts to train a statistical approach.
- The typical (normal) usage of other users was not believed to be a reasonable surrogate for intrusion attempts to train a statistical approach.
- Future intrusions need not follow historical methods.

This decision was probably the single most important and influential decision in shaping the statistical approach. If we had been willing to model intrusive behavior, then we could have applied the Neyman-Pearson Lemma (and the decades of statistical research that followed from this lemma) to develop an optimal statistical approach. The Neyman-Pearson Lemma (proven in the 1930s by two of the founders of modern statistical theory) states that if there are two known probability distributions, then the optimal statistical test for distinguishing between these two distributions is based on the ratio of their values. For example, let x denote an outcome (derived from audit records), let $n(x)$ denote the probability of that outcome for the host user, and let $m(x)$ denote the probability of that outcome for the intruder. Suppose that we want to determine the set of outcomes X such that X occurs very rarely for the host user (say no more than 0.1% of the time) but much more frequently for the intruder. The optimal selection for X is the set of values of x with the highest values for the ratio $m(x)/n(x)$. That is, we order all values for x according to the ratio of $m(x)/n(x)$ and add x values into X until $n(X)$ equals 0.1%. A restatement of this lemma in a Bayesian framework can be found in Helman et al. (1992).

If we had chosen to model intrusive behavior, then by access to the Neyman-Pearson Lemma we could have developed an optimal approach for detecting intrusive behavior. This would have allowed use of many standard statistical approaches (such as discriminant analysis, Hotelling's T^2 tests) that stemmed (either directly or indirectly) from this lemma. Instead, we were faced with the task of defining a region X without reference to assumed intruder behavior, but rather based on the fact that such behavior was unusual for the user (whether or not it would be unusual for the intruder).

4.2 Q: What Does the Space X of Outcomes Look Like? How Has the Unusual Aspects of this Space Influenced the NIDES Statistical Component?

Statistical intrusion approaches accept audit record inputs that define the outcomes of subject behavior. These outcomes define a space, which we have denoted as X . A member x belonging to X is the particular behavior that has been exhibited by a

subject. We want to decide whether x is representative of the normal behavior of the user or whether x is anomalous (and therefore indicative of misuse). For example, Helman et al. (1992) define two stationary stochastic processes on the outcome space — $N(x)$, which denotes normal or legitimate transactions, and $M(x)$, which denotes misuse transactions — that map values of x into probabilities or probability densities, and then later show the optimal test statistic for determining whether the particular value of x that was seen came from $N()$ or $M()$.

The ease in theoretically specifying that there is a space X belies the practical difficulties in actually defining X . (This difficulty directly affects the feasibility of implementing the Helman approach and has also influenced the NIDES statistical approach.) The following problems arise:

- Individual audit records typically do not contain complete values x within X . We observe a sequence of audit records, each of which contains limited information about user activity. Let us denote these audit records as r_t where t is the unique time stamp on the audit record. For example the value for r_t might denote that at the time that the last audit record was generated (i.e., time t) a particular command was being executed, a particular file was being accessed, the user was logged on from a particular terminal, and so forth. One method for specifying a value x is to let x be the complete past history of audit records or, perhaps more reasonably, all audit records generated over the last 10 minutes. That is, $x = \{r_t, r_{t-1}, r_{t-2}, r_{t-3}, r_{t-4}\}$, etc.
- A substantial number of the values of r_t are missing. For example, a particular audit record might not contain information about files being used at that moment, or about I/O, memory, or CPU usage. In addition, because a user can (intentionally or unknowingly) spawn a number of processes — all of which can be simultaneously using computer resources (even on a single machine) — and each audit record is tied to a specific process, each audit record gives only partial information about the totality of activity that is going on in a single user's account at the time of audit record generation.
- Because audit records from separate processes can be intermingled (as different processes swap in and out of memory or are actually simultaneously executed on different processors), user behavior as recorded on audit records can appear to be quite erratic. This type of behavior is not well explored in the literature (which usually assumes rather well-behaved time series or stochastic processes).
- Because there may be long periods without any audit activity and periods of time with very intense activity, it is not immediately apparent how many audit records should be included in x . For example, if x is defined in terms of clock time (say 5 minutes), there will be periods in which x is composed of thousands of audit records, and many other times when x is empty. On the other hand, if x is defined in terms of audit records (say a few hundred audit records), there will

be times when x denotes activity over a very short period of time (i.e., seconds or minutes) and times when x denotes activity over much longer periods of time (i.e., hours of near inactivity).

This discussion demonstrates the following:

- The dimensionality of x is potentially very large. At any particular time there may be many different processes executing (or temporarily in abeyance). Each individual process can have many different files open. Often, each process can have executed numerous commands. Whenever audit records address only very limited single actions (such as is the case with Sun-Unix audit records), we would want x to span a reasonable period of time. There is considerable uncertainty concerning whether to measure time by the clock or number of audit records, and either way can lead to undesirable results.
- The space X is at least as complicated as any outcome space considered in the statistical literature. The vast majority of statistical literature deals with very simple outcome spaces (e.g., independent or serially correlated outcomes, each of a fixed constant dimensionality) that cannot be applied to X without considerable modification.

In defining the NIDES statistical component we found it useful to let X be the space of all prior behavior (extending backwards in time to the first actions of the user). However, we weighted the audit records within x so that the more recent audit records were much more influential in our decision making. Furthermore, we found that by defining a set of statistics (i.e., long-term profiles) that were easily updatable exponentially weighted sums, we were able to avoid much of the complexity in the space X . For example, we could update our statistics only by looking at the most recent audit record, without having to consider the relationships of the current audit record to previous audit records.

4.3 Q: How Has the Size of the Space X Affected the NIDES Statistical Algorithm?

One tempting approach to anomaly detection is to enumerate all of the possible outcomes, observe the empirical probability for each of those outcomes for the normal user, and then declare a recent outcome (i.e., an audit record or set of audit records) to be anomalous if it is one of the outcomes that historically has rarely been seen. Unfortunately, the size of the space X precludes this approach.

The size of the space X is huge. Consider only two measures — command name used and file name used. A typical user might use 100 different commands and have a file system with 1000 different file names (excluding temporary files that are generated

by processes executed by the user). In this event, the space X has 100,000 possible values (although only a fraction of these possible values might actually have been experienced). Add measures for host used, time of day, CPU time used, I/O amount, memory usage, and so forth, and the potential size of the space X becomes enormous. Conservatively, the size of the space X would contain trillions of possible values. Even if we limit the space X to those values actually experienced by the user, X would contain millions or billions of values. Simply enumerating these values (so we can keep a record of their historical non-zero probabilities) might overwhelm our storage capability, and processing this list would unacceptably slow down our algorithm.

In addition to processing capability problems inherent in considering the full space of outcomes, there is an even greater problem concerning the sample size requirements. As noted by Helman et al., "the sample size mandated by the enormity of the event space $[X]$... might exceed memory limitations, or perhaps, the [normal and misuse] processes ..." cannot be assumed to be stationary over the time required to accumulate the requisite historical database." By the time we gather enough data to obtain reliable probabilities for the events in X , the user will have changed job assignments, changed jobs, or possibly died. Even if user behavior were extremely stable over time, we doubt that any system security officer would be willing to wait one or more years before starting anomaly detection on a new user.

On the basis of the above discussion, we conclude that *in general* the "tempting" procedure of enumerating all possible approaches is infeasible. However, there is one important exception. If the event space X can be reduced in size sufficiently, then the approach is implementable. This appears to be the case in the Wisdom and Sense algorithm. Because W&S relies on higher level auditing information (essentially information at the network level about logins) the size of the space X is severely restricted. Even in this restricted space, processing is a challenging task, and the W&S developers deserve praise for the clever algorithms that they use to process and encode the data. We decided not to pursue this method of restricting the scope of data to process statistically with NIDES for three reasons: (1) we felt that the expert system could do a fairly good job of detecting many types of unusual login activity, (2) we wanted NIDES to be capable of processing more detailed information about user activity, and (3) NIDES was designed to be extensible and allow the inclusion of modules such as W&S.

Our approach to the problems raised by the size of X was to assess the unusualness in each measure individually. (After we had assessed the unusualness of each measure individually, we summed these assessments to obtain an overall assessment of unusualness.) This approach dramatically lowers storage and the sample size requirements. The corresponding loss is that we cannot distinguish changes in the combinations of values of measures. For example, if a user only executes command "a" with file "b" and command "c" with file "d", then an intruder onto this user's account who executes command "a" with file "d" and command "c" with file "b" will not be flagged as anomalous. (This limitation is not inherent in the NIDES algorithm. Although

not currently part of NIDES, the NIDES statistical algorithm could be rewritten so that any two measures that are currently considered separately could also be considered jointly. For example, command name and file name could be jointly considered as a single measure, which takes categorical measures that are specified by both the command name and the file name.)

4.4 Q: What is the Relationship of the NIDES Statistical Algorithm to the Work of Helman et al.?

As mentioned earlier, the approach of Helman et al. is essentially to derive an optimal test statistic under the assumption that the probability distribution for misuse behavior is known. Helman et al. realize, however, that in practice probability distributions for misuse behavior will probably not be known. They therefore discuss two assumptions that can be made that dramatically simplify forms for these probability distributions:

- Independence — Under this assumption each measure (denoted by Helman as an attribute) is independent of each other attribute.
- Uniformity — Under this assumption, misuse behavior for each measure is assumed to be uniformly distributed over the possible values for the measure.

Because the NIDES statistical component does not assume a known probability distribution for misuse behavior, it cannot be easily related to Helman's approach. But we can make some observations:

- We do not believe that assumptions of independence or uniformity of misuse behavior are justified. Optimal procedures for detecting independent and uniform misuse behavior might be very poor in detecting actual misuse behavior. Nevertheless, it is desirable, in the absence of knowledge about misuse behavior, that an anomaly detector be at least capable of detecting independent and uniform misuse behavior.
- Because the NIDES statistical component uses a chi-square-like statistic to measure the differences between the short-term and long-term probability distributions for each individual measure, it will be capable of (and probably fairly good at) detecting uniform and independent misuse behavior.
- The NIDES statistical component can detect misuse behavior that is neither uniform nor independent. This is particularly true with respect to nonuniformity. Chi-square-type statistics are good at detecting an extremely broad range

of probability distributions that are not like the normal probability distribution. Their principal deficiency is that they are not as powerful as statistical tests formulated on the basis of a priori knowledge of the specific probability distribution of misuse behavior. That is, the chi-square-type statistics spread their power over a broad (in fact, infinite) range of probability distributions of misuse behavior.

- The NIDES statistical component does not assume independence of measures under either the normal or misuse behavior, nor does its validity depend on the independence of measures (that is, its false positive and false negative rates are uninfluenced by the dependence of measures). However, because the NIDES algorithm as currently configured does not place great emphasis on the correlation among measures, it is quite compatible with the assumption of independence of misuse behavior. In combination with the comments earlier about the uniformity assumption, we would expect the NIDES statistical algorithm to perform well when the misuse behavior is specified to be uniform and independent, although it will not perform as well as the optimal statistical test.
- Another type of misuse behavior that has not been explicitly considered by Helman is when the misuse behavior does not affect the marginal distributions of any measure, but does affect the correlation among the measures. As currently configured, the NIDES statistical component will not be particularly powerful in detecting this type of misuse behavior (unless the direction of change is to increase the correlation coefficient among the measures). However, it is possible to modify the algorithm in such a way that correlations are explicitly considered, thus, considerably strengthening the ability to detect changes in correlations. We view this as a potentially valuable area of future research.

4.5 Q: Can the NIDES Statistical Algorithm Consider Pairs of Measures Simultaneously (for example, command name and file name used)?

Theoretically, the NIDES statistical component can consider pairs of measures simultaneously, although this requires defining a new measure that is the cross product of two other measures. For example, consider the measure of name of file accessed and the measure of time of day. A composite measure might be defined as the two-tuple consisting of the name of the file accessed and the time of day of the access (perhaps categorized by hour of day). Each value for the two-tuple is a category for the new composite measure. The composite measure could thus, be handled by the NIDES statistical measure as a categorical measure in essentially the same way as

any other categorical measure. There is no need to limit this consideration to pairs; this approach would apply equally well for triplets, quadruplets, and so forth.

The theoretical advantage of considering pairs of measures (or equivalently, of constructing new measures that are two-tuples of existing measures) is that interactions between measures can be examined. Continuing our example, suppose that user A (who works 8 a.m. to 5 p.m. on weekdays) accesses file B daily, but only during the hour after lunch. One day, however, file B is used frequently between 9 a.m. and 10 a.m. The NIDES approach as currently implemented would not detect such usage as anomalous, because it is not unusual for user A to use file B, nor is it unusual for this user to access the computer during these hours. The use of a two-tuple approach would be necessary for NIDES to detect as anomalous the use of file B during the morning hours.

We have not added a mechanism in the NIDES computer program for constructing new measures from groups of other measures. The primary reason for not doing so is that we have concerns that the number of categories in composite measures would become so large, and the frequency with which each category is seen would be so small, that the NIDES methodology would cease to function well. The NIDES methodology works best when user behavior is concentrated in relatively few categories.

4.6 Q: Why Does the NIDES Statistical Algorithm Use Exponentially Weighted Sums?

The NIDES statistical algorithm uses exponentially weighted sums, rather than a window of past values, for three primary reasons. First, we want to reduce storage requirements and accelerate processing. Use of a window of past values would require that we store all the past values in the window. Each time a new audit record of the appropriate type arrived, we would have to recall those values, delete the oldest value, and add the newest value to the list. We would then recalculate our test statistic to accommodate the deletion of the oldest value and the addition of the newest value. This is considerably more effort and requires more storage than multiplying our summary weighted sum statistic by a decay constant and adding in the value of the newest audit record. Second, we use an exponentially weighted sum because we do not have any "natural" boundaries that delineate past from current behavior and therefore would not know how long the window should extend into the past. Third, we do want more distant behavior to count less towards anomalies than more recent behavior. This does not occur with a window approach (except that audit records further in the past than the window boundary do not count towards anomaly detection), but occurs smoothly and gradually with an exponentially weighted sum approach.

4.7 Q: Why Does the NIDES Statistical Algorithm Treat Counting Measures as if They Were Categorical Measures?

The NIDES statistical algorithm converts counting measures to categorical measures by defining categories that are ranges of counts. We do so because we wanted to be very general with respect to the types of changes in counting distributions that we would detect as anomalous. For example, consider a user who has two modes of using the computer — reading mail and running large spreadsheets. The first type of activity would result in low CPU usage, and the second type of activity would result in high CPU usage. Therefore, this user's CPU usage would tend to be bimodal, with very few audit records recording intermediate CPU values. In such a circumstance, if we observed a large number of audit records that reflected activity with an average amount of CPU usage (that is, audit records in the rarely occurring center bins of the CPU usage distribution) we would tend to be suspicious. Such activity could not be detected if we only examined statistics such as mean CPU usage. Rather, it would be necessary to examine the entire distribution of CPU usage. We found that the most convenient method for examining the entire distribution for counting measures was to redefine them as categorical measures. This allowed us to use the already-existing methodology for categorical measures.

Other statistical methods for examining the entire distribution of counting measures have been developed in the statistical literature. For example, Kolmogorov statistics allow the comparison of an empirical cumulative distribution with a long-term cumulative distribution. However, such procedures tend to either be computationally intensive or require substantial amounts of storage, and few have been adapted

4.8 Q: Why Does the NIDES Statistical Algorithm Use Exponential Weighting Based on Counts of Audit Records (rather than Clock Time) for Nearly All of the Measures?

The NIDES statistical algorithm differentiates between the volume of audit records that are currently being received, the general type of activity that is represented in those audit records, and the specific type of behavior that is occurring within each measure. For the purpose of monitoring the volume of audit records we use three audit record intensity measures, which currently have half-lives of 1, 10, and 60 minutes. For these three measures the exponential weighting is based on time. A single measure, denoted audit record distribution, tracks what percentage of the

most recent records have “triggered” or “touched” different measures. This provides an overview of the general types of activity that have recently taken place. For example, the audit record distribution measure would examine the percentage of recent audit records that concern a file access (that is, touch the file access measure). The audit record distribution measure uses exponential weighting based on counts of audit records. The remaining measures track the type of behavior that has transpired within a measure. For example, the file access measure examines which specific files were recently accessed, and whether or not these were unusual files.

In previous versions of NIDES, we did not subdivide the volume of activity from the specific behavior within that activity. This resulted in an inability to distinguish a relatively low volume of highly suspicious behavior from a high volume of behavior with typical levels of suspiciousness. The heuristic solution was to “normalize” our measures by the volume of activity; this was functionally equivalent to weighting on the basis of audit record counts and separately tracking audit record volume. This change to weighting on the basis of audit record count also facilitated the construction of short-term and long-term profiles, and the use of the chi-square statistic to compare those profiles. (Exponential weighting based on time is not compatible with the chi-square approach to profile comparison, and, as a result, the three intensity measures use a different approach.) There is a secondary advantage to the separation of audit record volume from the type of activity in those records — the measures are easier to interpret. We can separately detect whether the volume of activity is too high, whether the general type of activity within that volume is unusual (and in what ways it is unusual), and whether specific behavior within the activity is unusual (and in what ways it is unusual).

4.9 Q: For What Types of Computer Systems Will the NIDES Statistical Component be Most Applicable? Least Applicable?

The NIDES statistical component was developed to be very general and applicable across a wide range of systems and types of “subjects.” Thus, it is our belief that the NIDES statistical component should perform well in many circumstances. However, there is no way of knowing exactly how well the statistical component will perform without actually testing it.

It is important to distinguish between relative and absolute performance. Some circumstances are generally favorable to statistical anomaly detection, including stable subject profiles, users with sharply defined and restricted behavior, and substantial differences in misuse and normal behavior. Statistical systems, including NIDES, will tend to perform better in these circumstances. Conversely, when these circumstances are missing, all statistical systems will tend to yield poor results. We have no way of knowing how well the NIDES statistical component performs relative to other sta-

Answers

45

tistical approaches because other approaches have not been available to us to test until relatively recently, nor have we been funded to perform such testing. We believe that such testing could be very beneficial to statistical algorithm development, not so much to determine which statistical algorithms are better, but rather to contribute to an understanding of why they perform well or not, so that we can use such knowledge to improve them further.

Because our most important sources of audit record data were our own Sun-Unix systems and because the immediate client's systems were also Sun-Unix, there are undoubtedly aspects of our statistical component that are rather "tailored" to these types of computer systems and users. We have also had an opportunity to test our system in two other very different environments. One of these environments was IBM mainframe audit record data gathered in field offices of a federal agency, where the predominant usage was among clerks entering and retrieving data from a fixed number of large databases. The second use was data where the subjects were applications, and the intent of using the NIDES statistical component was to differentiate among applications to detect masquerading by more resource intensive applications (constituting misuse of exported computers). In all three environments, the NIDES statistical system has performed very well. The last environment (which was known as "Safeguard" because the intent was to safeguard the usage of exported computers) is particularly noteworthy because it was almost diametrically opposed (in many important statistical ways) to the SRI Sun-Unix environment in which NIDES was developed:

- The subjects in the SRI Sun-Unix environment are sophisticated computer users, while the subjects in Safeguard are applications.
- The volume of audit records in the SRI Sun-Unix environment averages about 2,000 per user-day; the volume in the Safeguard environment averages only about 10 per day.
- Individual audit records in the SRI Sun-Unix environment contain only a small fragment of the information necessary to determine the behavior of a user; each individual audit record in the Safeguard environment contains complete information about an invocation of an application.
- There is no natural beginning and ending to an SRI Sun-Unix user's activity, whereas in the Safeguard environment each audit record represents a complete invocation of an application and is the natural analysis unit.
- In the SRI Sun-Unix environment the most distinguishing measures tend to be those which are categorical, and therefore the NIDES statistical approach was developed with categorical measures in mind; in the Safeguard environment all measures are counting measures.

SYM_P_0078992

We were pleasantly surprised to find that the NIDES statistical component performed as well as it did in the Safeguard environment, given its stark differences with the SRI Sun-Unix environment, and the relatively small changes in the algorithm parameters that were implemented. This gives us reason to be optimistic that the NIDES statistical component can be applied in other circumstances, which were not in our minds when the algorithm was developed. However, as mentioned earlier, there is no substitute for actually testing the algorithm in those environments.

Bibliography

- [1] Michael R. Anderberg. *Cluster Analysis for Applications*. Academic Press, San Francisco, California, 1973.
- [2] Richard O. Duda and Peter E. Hart. *Pattern Recognition*. John Wiley and Sons, New York, 1973.
- [3] R. Jagannathan and Teresa Lunt. System design document: Next generation intrusion detection expert system (NIDES). SRI report, SRI International, Menlo Park, California, March 9, 1993.
- [4] T. F. Lunt, Ann Tamaru, Fred Gilham, R. Jagannathan, Caveh Jalali, H. S. Javitz, A. Valdes, P. G. Neumann, and T. D. Garvey. A Real-Time Intrusion-Detection Expert System (IDES), Final Technical Report. Technical report, Computer Science Laboratory, SRI International, Menlo Park, California, February 1992.
- [5] S. Karlin. *A First Course in Stochastic Processes*. Academic Press, New York, 1969.
- [6] R.A. Olshen L. Breiman, J.H. Friedman and C.J. Stone. *Classification and Regression Trees*. Wadsworth and Brooks/Cole Advanced Books and Software, Monterey, California, 1984.
- [7] *Wisdom and Sense Guidebook*. Los Alamos National Laboratory, Los Alamos, New Mexico. Undated.
- [8] Donald F. Morrison. *Multivariate Statistical Methods*. McGraw-Hill, San Francisco, California, 1976.
- [9] G. Liepins P. Helman and W. Richards. Foundations of intrusion detection. In *Proceedings of the IEEE Symposium on Computer Security*, 1992.
- [10] G. Shafer and editors J. Pearl. Readings in uncertain reasoning. Morgan Kaufmann Publishers, Inc., San Mateo, California, 1990.
- [11] S. Smaha. Haystack audit trail analysis system. Status Report HS_STAT.TXT, Haystack Laboratories, Colorado, August 15, 1990.